arXiv:1408.0424v1 [math.ST] 2 Aug 2014

Equivariant minimax dominators of the MLE in the array normal model

David Gerard[1] and Peter Hoff[1,2]

Departments of Statistics[1] and Biostatistics[2]

University of Washington

August 5, 2014

Abstract

Inference about dependencies in a multiway data array can be made using the array normal model, which corresponds to the class of multivariate normal distributions with separable covariance matrices. Maximum likelihood and Bayesian methods for inference in the array normal model have appeared in the literature, but there have not been any results concerning the optimality properties of such estimators. In this article, we obtain results for the array normal model that are analogous to some classical results concerning covariance estimation for the multivariate normal model. We show that under a lower triangular product group, a uniformly minimum risk equivariant estimator (UMREE) can be obtained via a generalized Bayes procedure. Although this UMREE is minimax and dominates the MLE, it can be improved upon via an orthogonally equivariant modification. Numerical comparisons of the risks of these estimators show that the equivariant estimators can have substantially lower risks than the MLE.

Keywords: Bayesian estimation, covariance estimation, Gibbs sampling, Stein's loss, tensor data

1 Introduction

The analysis of array-valued data, or tensor data, is of interest to numerous fields, including psychometrics [Kiers and Mechelen, 2001], chemometrics [Smilde et al., 2005, Bro, 2006], imaging [Vasilescu and Terzopoulos, 2003], signal processing [Cichocki et al., 2014] and machine learning

Email: gerard2@uw.edu, pdhoff@uw.edu. This research was partially supported by NI-CHD grant R01HD067509.

[Tao et al., 2005], among others [Kroonenberg, 2008, Kolda and Bader, 2009]. Such data consist of measurements indexed by multiple categorical factors. For example, multivariate measurements on experimental units over time may be represented by a three-way array $X = \{x_{i,j,t}\} \in \mathbb{R}^{m \times p \times t}$, with i indexing units, j indexing variables and t indexing time. Another example is multivariate relational data, where $x_{i,j,k}$ is the type-k relationship between person i and person j.

Statistical analysis of such data often proceeds by fitting a model such as $X = \Theta + E$, where Θ is low-dimensional and E represents additive residual variation about Θ. Standard models for Θ include regression models, additive effects models (such as those estimated by ANOVA decompositions) and unconstrained mean models if replicate observations are available. Another popular approach is to model Θ as being a low-rank array. For such models, ordinary least-squares estimates of Θ can be obtained via various types of tensor decompositions, depending on the definition of rank being used [De Lathauwer et al., 2000a,b, De Silva and Lim, 2008].

Less attention has been given to the analysis of the residual variation E. However, estimating and accounting for such variation is critical for a variety of inferential tasks, such as prediction, model-checking, construction of confidence intervals, and improved parameter estimation over ordinary least squares. One model for variation among the entries of an array is the array normal model [Akdemir and Gupta, 2011, Hoff, 2011] which is an extension of the matrix normal model [Srivastava and Khatri, 1979, Dawid, 1981], often used in the analysis of spatial and temporal data [Mardia, 1993, Shitan and Brockwell, 1995, Fuentes, 2006]. The array normal model is a class of normal distributions that are generated by a multilinear operator known as the Tucker product: A random K-way array X taking values in $\mathbb{R}^{p_1 \times \cdots \times p_K}$ has an array normal distribution if $X \stackrel{d}{=} \Theta + Z \times \{A_1, \ldots, A_K\}$, where "$\times$" denotes the Tucker product (described further in Section 2), Z is a random array in $\mathbb{R}^{p_1 \times \cdots \times p_K}$ having i.i.d. standard normal entries, and A_k is a $p_k \times p_k$ nonsingular matrix for each $k \in \{1, \ldots, K\}$. Letting $\Sigma_k = A_k A_k^T$ and "$\otimes$" denote the Kronecker product, we write

$$X \sim N_{p_1 \times \cdots \times p_K}(\Theta, \Sigma_K \otimes \cdots \otimes \Sigma_1). \tag{1}$$

A maximum likelihood estimate (MLE) for the parameters in (1) can be obtained via an iterative coordinate descent algorithm [Hoff, 2011], which is a generalization of the iterative "flip-flop" algorithm developed in Mardia [1993] and Dutilleul [1999], or alternatively the optimization procedures described in Wiesel [2012a]. However, based on results for the multivariate normal model,

one might suspect that the MLE lacks desirable optimality properties: In the multivariate normal model, James and Stein [1961] showed that the MLE of the covariance matrix is neither admissible nor minimax. This was accomplished by identifying a minimax and uniformly optimal equivariant estimator that is different from the (equivariant) MLE, and therefore dominates the MLE. As pointed out by James and Stein, this equivariant estimator is itself inadmissible, and improvements to this estimator have been developed and studied by Stein [1975], Takemura [1983], Lin and Perlman [1985], and Haff [1991], among others.

This article develops similar results for the array normal model. In particular, we obtain a procedure to obtain the uniformly minimum risk equivariant estimator (UMREE) under a lower-triangular product group of transformations for which the model (1) is invariant. Unlike for the multivariate normal model, there is no simple characterization of this class of equivariant estimators. However, results of Zidek [1969] and Eaton [1989] can be used to show that the UMREE can be obtained from the Bayes decision rule under an improper prior, which we derive in Section 2. In Section 3 we obtain the posterior distribution under this prior, and show how it can be simulated from using a Markov Chain Monte Carlo (MCMC) algorithm. Specifically, the MCMC algorithm is a Gibbs sampler that involves simulation from a class of distributions over covariance matrices, which we call the "mirror-Wishart" distributions.

In Section 4.1 we develop a version of Stein's loss function for covariance estimation in the array normal model, and show how the Gibbs sampler of Section 3 can be used to obtain the UMREE for this loss. We discuss an orthogonally equivariant improvement to the UMREE in Section 4.2, which can be seen as analogous to the estimator studied by Takemura [1983]. Section 4.3 compares the risks of the MLE, UMREE and the orthogonally equivariant estimator as a function of the dimension of X in a small simulation study. A discussion follows in Section 5. Proofs are contained in an appendix.

2 An invariant measure for the array normal model

2.1 The array normal model

The array normal model on $\mathbb{R}^{p_1 \times \cdots \times p_K}$ consists of the distributions of random K-arrays $X \in \mathbb{R}^{p_1 \times \cdots \times p_K}$ for which

$$X \stackrel{d}{=} \Theta + Z \times \{A_1, \ldots, A_K\} \tag{2}$$

for some $\Theta \in \mathbb{R}^{p_1 \times \cdots \times p_K}$, nonsingular matrices $A_k \in \mathbb{R}^{p_k \times p_k}, k = 1, \ldots, K$ and a random $p_1 \times \cdots \times p_K$ array Z with i.i.d. standard normal entries. Here, "$\times$" denotes the *Tucker product*, which is defined by the identity

$$\text{vec}(Z \times \{A_1, \ldots, A_K\}) = (A_K \otimes \cdots \otimes A_1)z, \tag{3}$$

where "$\otimes$" is the Kronecker product and $z = \text{vec}(Z)$, the vectorization of Z. This identity can be used to find the covariance of the elements of a random array satisfying (2): Letting x, z, θ be the vectorizations of X, Z, Θ, we have

$$\begin{aligned}
\text{Cov}[x] &= \text{E}[(x-\theta)(x-\theta)^T] \\
&= \text{E}[(A_K \otimes \cdots \otimes A_1) z z^T (A_K^T \otimes \cdots \otimes A_1^T)] \\
&= (A_K \otimes \cdots \otimes A_1)(A_K^T \otimes \cdots \otimes A_1^T) = (A_K A_K^T \otimes \cdots \otimes A_1 A_1^T),
\end{aligned}$$

and so the array normal distributions correspond to the multivariate normal distributions with separable (Kronecker structured) covariance matrices.

A useful operation related to the Tucker product is the *matricization* operation, which reshapes an array into a matrix along an index set, or *mode*. For example, the *mode-k matricization* of Z is the $p_k \times (\prod_{l:l \neq k} p_l)$-dimensional matrix $Z_{(k)}$ having rows equal to the vectorizations of the "slices" of Z along the kth index set. An important identity involving the Tucker product is that if $Y = Z \times \{A_1, \ldots, A_K\}$ then

$$Y_{(k)} = A_k Z_{(k)} \left(A_K^T \otimes \cdots \otimes A_{k+1}^T \otimes A_{k-1}^T \otimes \cdots \otimes A_1^T\right). \tag{4}$$

As shown in Hoff [2011], a direct application of this identity gives

$$E\left[(X_{(k)} - \Theta_{(k)})(X_{(k)} - \Theta_{(k)})^T\right] = c_k A_k A_k^T,$$

where c_k is a scalar. This shows that $A_k A_k^T$ can be interpreted as the covariance among the p_k slices of the array X along its kth mode.

The array normal model can be parameterized in terms of a mean array $\mathrm{E}[X] = \Theta \in \mathbb{R}^{p_1 \times \cdots \times p_K}$ and covariance $\mathrm{Cov}[\mathrm{vec}(X)] = \sigma^2(\Sigma_K \otimes \cdots \otimes \Sigma_1)$, where $\sigma^2 > 0$ and for each k, $\Sigma_k \in \mathcal{S}_{p_k}^+$, the set of $p_k \times p_k$ positive definite matrices. To make the parameterization identifiable, we restrict the determinant of each Σ_k to be one. Denote by $\mathcal{S}_{\mathbf{p}}^+$ this parameter space, that is, the values of $(\sigma^2, \Sigma_1, \ldots, \Sigma_K)$ for which $|\Sigma_k| = 1$, $k = 1, \ldots, K$. Under this parameterization, we write $X \sim N_{p_1 \times \cdots \times p_K}(\Theta, \sigma^2(\Sigma_K \otimes \cdots \otimes \Sigma_1))$ if and only if $X \stackrel{d}{=} \Theta + \sigma Z \times \{\Psi_1, \ldots, \Psi_K\}$, where for each k, Ψ_k is a matrix such that $\Psi_k \Psi_k^T = \Sigma_k$.

Given a sample $X_1, \ldots, X_n \sim$ i.i.d. $N_{p_1 \times \cdots \times p_K}(\Theta, \sigma^2(\Sigma_K \otimes \cdots \otimes \Sigma_1))$, the $(K+1)$-array X obtained by "stacking" $X_1, \ldots, X_n$ along a $(K+1)$st mode also has an array normal distribution,

$$X \sim N_{p_1 \times \cdots \times p_K \times n}\left(\Theta \circ \mathbf{1}_n, \sigma^2(I_n \otimes \Sigma_K \otimes \cdots \otimes \Sigma_1)\right),$$

where $\mathbf{1}_n$ is the $n \times 1$ vector of ones and "$\circ$" denotes the outer product. If $n > 1$ then covariance estimation for the array normal model can be reduced to the case that $\Theta = 0$. To see this, let H be a $(n-1) \times n$ matrix such that $HH^T = I_{n-1}$ and $H1_n = 0$. This implies that $H^T H = I_n - \mathbf{1}_n \mathbf{1}_n^T / n$. Letting $Y = X \times \{I_{p_1}, \ldots, I_{p_K}, H\}$, and $Y_{(K+1)}$ be the mode-$(K+1)$ matricization of Y, we have

$$\mathrm{E}[Y_{(K+1)}] = H\mathrm{E}[X_{(K+1)}] = H\mathbf{1}_n \mathrm{vec}(\Theta)^T = \mathbf{0},$$

and so Y is mean-zero. Using identity (3), the covariance of $\mathrm{vec}(Y)$ can be shown to be $\sigma^2(HH^T \otimes \Sigma_K \otimes \cdots \otimes \Sigma_1) = \sigma^2(I_{n-1} \otimes \Sigma_K \otimes \cdots \otimes \Sigma_1)$, and so $Y \sim N_{p_1 \times \cdots \times p_K \times (n-1)}(0, \sigma^2(I_{n-1} \otimes \Sigma_K \otimes \cdots \otimes \Sigma_1))$. For the remainder of this paper, we consider covariance estimation in the case that $\Theta = 0$.

2.2 Model invariance and a right invariant measure

Consider the model for an i.i.d. sample of size n from a p-variate mean-zero multivariate normal distribution, $X \sim N_{p \times n}(0, I_n \otimes \Sigma)$, $\Sigma \in \mathcal{S}_p^+$. Recall that $AX \sim N_{p \times n}(0, I_n \otimes A\Sigma A^T)$ for nonsingular matrices A, and so in particular this model is invariant under left multiplication of X by elements of G_p^+, the group of lower triangular matrices with positive diagonals. An estimator $\hat{\Sigma}$ mapping the sample space $\mathbb{R}^{p \times n}$ to $\mathcal{S}_p^+$ is said to be equivariant under this group if $\hat{\Sigma}(AX) = A\hat{\Sigma}(X)A^T$ for all $A \in G_p^+$ and $X \in \mathbb{R}^{p \times n}$. James and Stein [1961] characterized the class of equivariant estimators for this model, identified the UMREE under a particular loss function and showed that the UMREE

is minimax. Additionally, as the MLE XX^T/n is equivariant and different from the UMREE, the MLE is dominated by the UMREE.

We pursue analogous results for the array normal model by first reparameterizing in terms of the parameter $\Sigma^{1/2} = (\sigma, \Psi_1, \ldots, \Psi_K)$, so

$$X \sim N_{p_1 \times \cdots \times p_K \times n}\left(0, \sigma^2\left(I_n \otimes \Psi_K \Psi_K^T \otimes \cdots \otimes \Psi_1 \Psi_1^T\right)\right), \tag{5}$$

where $\sigma > 0$ and each Ψ_k is in the set $\mathcal{G}_{p_k}^+$ of $p_k \times p_k$ lower triangular matrices with positive diagonals and determinant 1. In this parameterization, Ψ_k is the lower triangular Cholesky square root of the mode-k covariance matrix Σ_k described in Section 2.1.

Define the group $\mathcal{G}_{\mathbf{p}}^+$ as

$$\mathcal{G}_{\mathbf{p}}^+ = \left\{A = (a, A_1, \ldots, A_K) \ : \ a > 0, A_k \in \mathcal{G}_{p_k}^+ \text{ for } k = 1, \ldots, K\right\},$$

where the group operation is

$$AT = (a, A_1, \ldots, A_K)(t, T_1, \ldots, T_K) = (at, A_1T_1, \ldots, A_KT_K).$$

Note that $\mathcal{G}_{\mathbf{p}}^+$ consists of the same set as the parameter space for the model, as parameterized in (5). If the group $\mathcal{G}_{\mathbf{p}}^+$ acts on the sample space by

$$g : X \mapsto aX \times \{A_1, \ldots, A_K, I_n\},$$

then as shown in Hoff [2011] it acts on the parameter space by

$$g : (\sigma, \Psi_1, \ldots, \Psi_K) \mapsto (a\sigma, A_1\Psi_1, \ldots, A_K\Psi_K),$$

which we write concisely as $g : \Sigma^{1/2} \mapsto A\Sigma^{1/2}$. An estimator, $\hat{\Sigma}^{1/2} = (\hat{\sigma}, \hat{\Psi}_1, \ldots, \hat{\Psi}_K)$, mapping the sample space $\mathbb{R}^{p_1 \times \cdots \times p_K \times n}$ to the parameter space $\mathcal{G}_{\mathbf{p}}^+$ is equivariant if

$$\hat{\Sigma}^{1/2}\left(aX \times \{A_1, \ldots, A_K, I_n\}\right) = (a, A_1, \ldots, A_K)\hat{\Sigma}^{1/2}\left(X\right).$$

For example, if $\hat{\Psi}_k$ is the estimator of Ψ_k when observing X, then $A_k\hat{\Psi}_k$ is the estimator when observing $aX \times \{A_1 \ldots, A_K, I_n\}$.

Unlike the case for the multivariate normal model, the class of $\mathcal{G}_{\mathbf{p}}^+$- equivariant estimators for the array normal model is not easy to characterize beyond the definition given above. However, in cases like the present one where the group space and parameter space are the same, the UMREE

under an invariant loss can be obtained as the generalized Bayes decision rule under a (generally improper) prior obtained from a right invariant (Haar) measure over the group [Zidek, 1969, Eaton, 1989]. The first step towards obtaining the UMREE is then to obtain a right invariant measure and corresponding prior. To do this, we first need to define an appropriate measure space for the elements of $\mathcal{G}_{\mathbf{p}}^{+}$. Recall that matrices A_k in $\mathcal{G}_{p_k}^{+}$ have determinant 1, and so one of the nonzero elements of A_k can be expressed as a function of the others. For the rest of this section and the next, we parameterize $A_k \in \mathcal{G}_{p_k}^{+}$ in terms of the elements $\{A_{k[i,j]} : 2 \leq i \leq p_k, 1 \leq j \leq i\}$, and express the upper-left element $A_{k[1,1]}$ as a function of the other diagonal elements, so that $A_{k[1,1]} = \prod_{i=2}^{p_k} (A_{k[i,i]})^{-1}$. The "free" elements of $A_k \in \mathcal{G}_{p_k}^{+}$ therefore take values in the space $\mathcal{A}_{p_k} = \{a_{i,i} > 0, a_{i,j} \in \mathbb{R} : 2 \leq i \leq p_k, 1 \leq j < i\}$.

Theorem 1. *A right invariant measure over the group $\mathcal{G}_{\mathbf{p}}^{+}$ is*

$$d\nu_r\left(a, A_1, \ldots, A_K\right) = \frac{1}{a}\left(\prod_{k=1}^{K}\prod_{i=2}^{p_k} A_{k[i,i]}^{i-2}\right)\, d\mu\left(a, A_1, \ldots, A_K\right),$$

where $d\mu$ is Lebesgue measure over $\mathbb{R}^{+} \times \mathcal{A}_{p_1} \times \cdots \times \mathcal{A}_{p_K}$.

We note that although the density given above is specific to the particular parameterization of the $\mathcal{G}_{p_k}^{+}$'s, the inference results that follow will hold for any parameterization.

Let $L : \mathcal{G}_{\mathbf{p}}^{+} \times \mathcal{G}_{\mathbf{p}}^{+} \to \mathbb{R}^{+}$ be an invariant loss function, so that $L(\Sigma^{1/2}, B) = L(A\Sigma^{1/2}, AB)$ for all A, B and $\Sigma^{1/2} \in \mathcal{G}_{\mathbf{p}_k}^{+}$. Theorem 6.5 of Eaton [1989] implies that the value of the UMREE when the array X is observed is the minimizer in $B = (b, B_1, \ldots, B_K)$ of the integral

$$\int_{\mathcal{G}_{\mathbf{p}}^{+}} L(A\Sigma_0^{1/2}, B) \times p(X|A\Sigma_0^{1/2})\; d\nu_r(A),$$

where $p(X|A\Sigma_0^{1/2})$ is the array normal density at the parameter value $A\Sigma_0^{1/2}$ and $\Sigma_0^{1/2}$ is an arbitrary element of $\mathcal{G}_{\mathbf{p}}^{+}$. Since the group action is transitive over the parameter space, and since the integral is right invariant, $\Sigma_0^{1/2}$ can be chosen to be equal to $(1, I_{p_1}, \ldots, I_{p_K})$. Furthermore, since the parameter space and group space are the same, replacing A with $\Sigma^{1/2}$ in the above integral indicates that the UMREE at X is the minimizer in B of

$$\int_{\mathcal{G}_{\mathbf{p}}^{+}} L(\Sigma^{1/2}, B) \times p(X|\Sigma^{1/2})\; d\nu_r(\Sigma^{1/2}),$$

that is, the UMREE is the Bayes estimator under the (improper) prior ν_r for $\Sigma^{1/2}$. This is summarized in the following corollary:

Corollary 1. *For an invariant loss function* $L : \mathcal{G}_{\mathbf{p}}^+ \times \mathcal{G}_{\mathbf{p}}^+ \to \mathbb{R}^+$ *the estimator* $\hat{\Sigma}^{1/2}$*, defined as*

$$\hat{\Sigma}^{1/2}(X) = \underset{B \in \mathcal{G}_{\mathbf{p}}^+}{\arg\min}\, E[L(\Sigma^{1/2}, B)|X], \tag{6}$$

uniformly minimizes the risk $E[L(\Sigma^{1/2}, \tilde{\Sigma}^{1/2}(X))|\Sigma^{1/2}]$ *among equivariant estimators* $\tilde{\Sigma}^{1/2}$ *of* $\Sigma^{1/2}$*. The expectation in (6) is with respect to the posterior density*

$$\begin{aligned} &p(\sigma, \Psi_1, \ldots, \Psi_K|X) \propto \\ &\sigma^{-np} \exp\left\{-\frac{1}{2\sigma^2}||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^2\right\} \frac{1}{\sigma} \prod_{k=1}^{K} \prod_{i=2}^{p_k} \Psi_{k[i,i]}^{i-2}, \end{aligned} \tag{7}$$

where $p = \prod_1^K p_k$*.*

In addition to uniformly minimizing the risk, the UMREE has two additional features. First, since any unique MLE is equivariant [Eaton, 1989, Theorem 3.2], the UMREE dominates any unique MLE, presuming the UMREE is not the MLE. Second, the UMREE under $\mathcal{G}_{\mathbf{p}}^+$ is minimax. This follows because $\mathcal{G}_{\mathbf{p}}^+$ is a subgroup of G_p^+, as $a(A_K \otimes \cdots \otimes A_1) \in G_p^+$ for all $a > 0$ and $A_k \in \mathcal{G}_{p_k}^+$. Since G_p^+ is a solvable group [James and Stein, 1961], this necessarily implies that $\mathcal{G}_{\mathbf{p}}^+$ is solvable [Rotman, 1995, Theorem 5.15]. By the results of Kiefer [1957] and Bondar and Milnes [1981], the equivariant estimator that minimizes (6) is minimax.

Note that because the prior ν_r is improper, the posterior (7) is not guaranteed to be proper. However, we are able to guarantee propriety if the sample size n is sufficiently large:

Theorem 2. *Let* $n > \prod_{k=1}^K p_k$*. For* $p(\sigma, \Psi_1, \ldots, \Psi_K|X)$ *defined in (7),*

$$\int_{\mathbb{R}^+ \times \mathcal{G}_{p_1}^+ \times \cdots \times \mathcal{G}_{p_K}^+} p(\sigma, \Psi_1, \ldots, \Psi_K|X) d\sigma d\Psi_1 \cdots d\Psi_K < \infty$$

The sample size in the Theorem is sufficient for propriety, but empirical evidence suggests that it is not necessary. For example, results from a simulation study in Section 4 suggest that, for some dimensions, a sample size of $n = 1$ is sufficient for posterior propriety and existence of an UMREE.

3 Posterior approximation

For the results in Section 2 to be of use, we must be able to actually minimize the posterior risk in Equation 6 under an invariant loss function of interest. In the next section, we will show that

the posterior risk minimizer under a multiway generalization of Stein's loss is given by posterior expectations of the form $\mathrm{E}[(\sigma^2\Sigma_k)^{-1}|X]$, where $\Sigma_k = \Psi_k\Psi_k^T$. Although these posterior expectations are not generally available in analytic form, they can be approximated using a MCMC algorithm. In this section, we show how a relatively simple Gibbs sampler can be used to simulate a Markov chain of values of $\Sigma^{1/2} = (\sigma, \Psi_1, \ldots, \Psi_K)$, having a stationary distribution equal to the desired posterior distribution given by Equation 7. These simulated values can be used to approximate the posterior distribution of $\Sigma^{1/2}$ given X, as well as any posterior expectation, in particular $\mathrm{E}[(\sigma^2\Sigma_k)^{-1}|X]$.

The Gibbs sampler proceeds by iteratively simulating values of $\{\sigma, \Psi_k\}$ from their full conditional distribution given the current values of $\{\Psi_1, \ldots, \Psi_{k-1}, \Psi_{k+1}, \ldots, \Psi_K\}$. This is done by simulating $\sigma^2\Sigma_k$ from its full conditional distribution, from which σ and Ψ_k can be recovered. One iteration of the Gibbs sampler proceeds as follows:

Iteratively for each $k \in \{1, \ldots, K\}$,

1. simulate $(\sigma^2\Sigma_k)^{-1} \sim \text{mirror-Wishart}_{p_k}(np/p_k, (X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T)^{-1})$;
2. set Ψ_k to be the lower triangular Cholesky square root of Σ_k.

In this algorithm, $X_{(k)} \in \mathbb{R}^{p_k \times np/p_k}$ is the mode-k matricization of X and $\Psi_{-k} = \Psi_K \otimes \cdots \otimes \Psi_{k+1} \otimes \Psi_{k-1} \otimes \cdots \otimes \Psi_1$. The mirror-Wishart distribution is a probability distribution on positive definite matrices, related to the Wishart distribution as follows:

Definition 1. *A random $q \times q$ positive definite matrix S has a mirror-Wishart distribution with degrees of freedom $\nu > 0$ and scale matrix $\Phi \in \mathcal{S}_q^+$ if*

$$S \stackrel{d}{=} UV^TVU^T,$$

where VV^T is the lower triangular Cholesky decomposition of a Wishart$_q(\nu, I_q)$-distributed random matrix and UU^T is the upper triangular Cholesky decomposition of Φ.

Some understanding of the mirror-Wishart distribution can be obtained from its expectation:

Lemma 1. *If $S \sim$ mirror-Wishart$_q(\nu, \Phi)$ then*

$$E[S] = \nu UDU^T$$

where UU^T is the upper triangular Cholesky decomposition of Φ and D is a diagonal matrix with entries $d_j = (\nu + q + 1 - 2j)/\nu,\ j = 1, \ldots, q$.

The calculation follows from Bartlett's decomposition, and is in the appendix. The implications of this for covariance estimation are best understood in the context of the multivariate normal model $X \sim N_{p\times n}(0, I_n \otimes \Sigma)$. In this case, for a given prior the Bayes estimator under Stein's loss is given by $\mathrm{E}[\Sigma^{-1}|X]^{-1}$ (see, for example Yang and Berger [1994]). Under Jeffreys' noninformative prior, $\Sigma^{-1} \sim \mathrm{Wishart}_p(n, (XX^T)^{-1})$ and so the Bayes estimator is XX^T/n. While unbiased, this estimator is generally thought of as not providing appropriate shrinkage of the sample eigenvalues. Note that under Jeffreys' prior, *a posteriori* we have $\Sigma^{-1} \stackrel{d}{=} UVV^TU^T$, where $VV^T \sim \mathrm{Wishart}_p(n, I_p)$ and UU^T is the upper triangular Cholesky decomposition of $(XX^T)^{-1}$. In contrast, under a right invariant measure as our prior we have $\Sigma^{-1} \stackrel{d}{=} UV^TVU^T$. The expectation of VV^T is nI, whereas the expectation of V^TV is nD, which provides a different pattern of shrinkage of the eigenvalues of XX^T. By Lemma 1 , the Bayes estimator under a right invariant measure as our prior in this case is given by $(nUDU^T)^{-1} = U^{-T}D^{-1}U^{-1}/n$, which is the UMREE obtained by James and Stein [1961]. Thus, the UMREE in the multivariate normal model corresponds to a Bayes estimator under a right invariant measure as our prior and mirror-Wishart posterior distribution.

The Gibbs sampler is based on the full conditional distribution of $(\sigma^2\Sigma_k)^{-1}$, which we derive from the full conditional density of $\{\sigma, \Psi_k\}$:

$$p(\sigma, \Psi_k) \propto |\sigma\Psi_k|^{-(np+1)/p_k} \exp\left\{-\mathrm{tr}\left((\sigma^2\Psi_k\Psi_k^T)^{-1}X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T\right)/2\right\} \prod_{i=2}^{p_k} \Psi_{k[i,i]}^{i-2},$$

where dependence of the density on $\{\Psi_1, \ldots, \Psi_{k-1}, \Psi_{k+1}, \ldots, \Psi_K, X\}$ has been made implicit. Now set $L_k = \sigma\Psi_k$. The full conditional density of L_k can be obtained from that of $\{\sigma, \Psi_k\}$ and the Jacobian of the transformation.

Lemma 2. *The Jacobian of the transformation $g(\sigma, \Psi_k) = \sigma\Psi_k$, mapping $\mathbb{R}^+ \times \mathcal{G}_{p_k}^+$ to $G_{p_k}^+$ is*

$$J(\sigma, \Psi_k) \propto \sigma^{p_k(p_k+1)/2-1}\Psi_{k[1,1]}.$$

Since $L_k = \sigma\Psi_k$, we have $\sigma = |L_k|^{1/p_k}$ and $\Psi_{k[i,i]} = L_{k[i,i]}/\sigma = L_{k[i,i]}/|L_k|^{1/p_k}$. Lemma 2 implies

$$\begin{aligned} p(L_k) \propto |L_k^T|^{-(np+1)/p_k} &\exp\left\{-\mathrm{tr}\left((L_kL_k^T)^{-1}X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T\right)/2\right\} \\ &\times \prod_{i=2}^{p_k}\left(L_{k[i,i]}/|L_k|^{1/p_k}\right)^{i-2}\left(|L_k|^{1/p_k}\right)^{-p_k(p_k+1)/2+1}\left(L_{k[1,1]}/|L_k|^{1/p_k}\right)^{-1}, \end{aligned}$$

which, through straightforward calculations, can be shown to be proportional to

$$\left(\prod_{i=1}^{p_k} L_{k[i,i]}^{i-np/p_k-p_k-1}\right)\exp\left\{-\mathrm{tr}\left((L_kL_k^T)^{-1}X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T\right)/2\right\}.$$

We now "absorb" $X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T$ into L_k. First, take the lower triangular Cholesky decomposition of $X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T = \Phi_k\Phi_k^T$ so that

$$\left(X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T\right)^{-1} = \Phi_k^{-T}\Phi_k^{-1}.$$

We have

$$\begin{aligned} p(L_k) &\propto \left(\prod_{i=1}^{p_k} L_{k[i,i]}^{i-np/p_k-p_k-1}\right) \exp\left\{-\text{tr}\left((L_kL_k^T)^{-1}\Phi_k\Phi_k^T\right)/2\right\} \\ &\propto \left(\prod_{i=1}^{p_k} L_{k[i,i]}^{i-np/p_k-p_k-1}\right) \exp\left\{-\text{tr}\left((\Phi_k^{-1}L_k(\Phi_k^{-1}L_k)^T)^{-1}\right)/2\right\}. \end{aligned}$$

Now let $W_k = \Phi_k^{-1}L_k$, so that $L_k = \Phi_k W_k$. This change of variables has Jacobian $J(W_k) = \prod_{i=1}^{p_k}\Phi_{k[i,i]}^i$ [Eaton, 1983, Proposition 5.13], so that

$$p(W_k) \propto \left(\prod_{i=1}^{p_k} W_{k[i,i]}^{i-np/p_k-p_k-1}\right) \exp\left\{-\text{tr}\left((W_kW_k^T)^{-1}\right)/2\right\}. \tag{8}$$

Note that the distribution of W_k does not depend on Ψ_{-k}. Now compare equation (8) to the density of the lower triangular Cholesky square root W of an inverse-Wishart distributed random matrix

$$WW^T \sim \text{inverse-Wishart}_{p_k}\left(np/p_k, I_{p_k}\right),$$

given by

$$p(W) \propto \left(\prod_{i=1}^{p_k} W_{[i,i]}^{-np/p_k-i}\right) \exp\left\{-\text{tr}((WW^T)^{-1})/2\right\}. \tag{9}$$

The conditional densities of the off-diagonal elements of W_k and W given the diagonal elements clearly have the same form. The diagonal elements of W_k and W in (8) and (9) turn out to be square roots of inverse-gamma distributed random variables, but with different shape parameters. To show this, we first derive the conditional densities of the off-diagonal elements of W:

Lemma 3. *(Bartlett's decomposition for the inverse-Wishart) Let W be the lower triangular Cholesky square root of an inverse-Wishart distributed matrix, so $WW^T \sim$ inverse-Wishart$_{p_k}(\nu, I_{p_k})$. Then for each $i = 1, \ldots, p_k$,*

$$\begin{aligned} &W_{[i,i]}^2 \sim \ \textit{inverse-gamma}([\nu - p_k + i]/2, 1/2), \ \textit{and} \\ &W_{[i,1:(i-1)]}|W_{[i,i]}, W_{[1:(i-1),1:(i-1)]} \\ &\sim N_{i-1}\left(0, W_{[i,i]}^2 W_{[1:(i-1),1:(i-1)]}^T W_{[1:(i-1),1:(i-1)]}\right). \end{aligned}$$

Here, $W_{[1:(i-1),1:(i-1)]}$ denotes the submatrix of W made up of the first $(i-1)$ rows and columns, and $W_{[i,1:(i-1)]}$ is the vector made up of the first $(i-1)$ elements of the ith row.

By Lemma 3, if $WW^T \sim$ inverse-Wishart$(np/p_k, I_{p_k})$ then the squared diagonal elements of W are independent inverse-gamma$((np/p_k - p_k + i)/2, 1/2)$ random variables. This tells us that

$$\int \exp\left\{-\mathrm{tr}((WW^T)^{-1})/2\right\} \prod_{i>j} dW_{[i,j]} \;\propto\; \prod_{i=1}^{p_k} W_{[i,i]}^{p_k-1} \exp\left\{-1/(2W_{[i,i]}^2)\right\}.$$

This result allows us to integrate (8) with respect to the off-diagonal elements of W_k, giving

$$\begin{aligned}
&\int \left(\prod_{i=1}^{p_k} W_{k[i,i]}^{i-np/p_k-p_k-1} \right) \exp\left\{-\mathrm{tr}\left((W_k W_k^T)^{-1}\right)/2\right\} \prod_{i>j} dW_{[i,j]} \\
&\propto W_{k[i,i]}^{i-np/p_k-2} \exp\left\{-1/(2W_{k[i,i]}^2)\right\}.
\end{aligned}$$

A change of variables implies that the $W_{k[i,i]}^2$'s are independent, and

$$W_{k[i,i]}^2 \sim \text{ inverse-gamma}([np/p_k - i + 1]/2, 1/2). \tag{10}$$

This completes the characterization of the distribution of W_k: The distribution of the diagonal elements is given by (10) and the conditional distribution of the off-diagonal elements given the diagonal can be obtained from Lemma 3. Finally, this distribution can be related to a Wishart distribution via the following lemma:

Lemma 4. *Let W_k be a random $p_k \times p_k$ lower triangular matrix such that*

$$\begin{aligned}
&W_{k[i,i]}^2 \sim \;\; \textit{inverse-gamma}\left([\nu - i + 1]/2, 1/2\right), \;\; \textit{and} \\
&W_{k[i,1:(i-1)]} | W_{k[1:(i-1),1:(i-1)]}, W_{k[i,i]} \\
&\sim N_{i-1}\left(0, W_{k[i,i]}^2 W_{k[1:(i-1),1:(i-1)]}^T W_{k[1:(i-1),1:(i-1)]}\right).
\end{aligned}$$

Then the elements of $V_k = W_k^{-1}$ are distributed independently as

$$\begin{aligned}
&V_{k[i,i]}^2 \sim \textit{gamma}([\nu - i + 1]/2, 1/2), \;\; i = 1, \ldots, q \\
&V_{k[i,j]} \sim N(0,1), \;\; i \neq j.
\end{aligned}$$

Note that the matrix V_k is distributed as the lower triangular Cholesky square root of a Wishart distributed random matrix. Applying the lemma to W_k, for which $\nu = np/p_k$, we have that $V_k = W_k^{-1} = (\Phi_k^{-1} L_k)^{-1} = L_k^{-1}\Phi_k = \frac{1}{\sigma}\Psi_k^{-1}\Phi_k$ is equal in distribution to the lower triangular

Cholesky square root of a Wishart$_{p_k}(np/p_k, I_{p_k})$ random matrix. That is, the precision matrix $(\sigma^2\Psi_k\Psi_k^T)^{-1} = \Psi_k^{-T}\Psi_k^{-1}/\sigma^2$ is conditionally distributed as

$$\frac{1}{\sigma^2}\Psi_k^{-T}\Psi_k^{-1}|\Psi_{-k} \stackrel{d}{=} \Phi_k^{-T}V^TV\Phi_k^{-1}, \text{ where}$$
$$VV^T \sim \text{Wishart}_{p_k}(np/p_k, I_{p_k}) \text{ and } \Phi_k\Phi_k^T = X_{(k)}\Psi_{-k}^{-T}\Psi_{-k}^{-1}X_{(k)}^T.$$

We say the matrix, $\Phi_k^{-T}V^TV\Phi_k^{-1}$ has a *mirror-Wishart* distribution because $\Phi_k^{-T}VV^T\Phi_k^{-1}$ would have a Wishart distribution. This completes the derivation of the full conditional distribution of $\sigma^2\Sigma_k = \sigma^2\Psi_k\Psi_k^T$.

Although not necessary for posterior approximation, the full conditional distribution of σ given $\Psi_1, \ldots, \Psi_K$ and X is easy to derive. The posterior density is

$$p(\sigma) \propto \sigma^{-(np+1)} \exp\left\{-\left|\left|X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}\right|\right|^2/(2\sigma^2)\right\}.$$

Letting $\gamma = 1/\sigma^2$, we have

$$p(\gamma) \propto \gamma^{np/2-1} \exp\left\{-\gamma\left|\left|X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}\right|\right|^2/2\right\},$$

and so the full conditional distribution of $1/\sigma^2$ is

$$\text{gamma}(np/2, ||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^2/2).$$

4 Estimation under multiway Stein's loss

4.1 The UMREE for multiway Stein's loss

A commonly used loss function for estimation of a covariance matrix Σ is Stein's loss,

$$L_S(S, \Sigma) = \text{tr}\left(S\Sigma^{-1}\right) - \log\left|S\Sigma^{-1}\right| - p, \quad \Sigma, S \in \mathcal{S}_p^+.$$

First introduced by James and Stein [1961], Stein's loss has been proposed as a reasonable and perhaps better alternative to quadratic loss for evaluating performance of covariance estimators. For example, Stein's loss, unlike quadratic loss, does not penalize overestimation of the variances more severely than underestimation.

Recall from Section 2 that the array normal model can be parameterized in terms of $\Sigma = (\sigma^2, \Sigma_1, \ldots, \Sigma_K) \in \mathcal{S}_{\mathbf{p}}^+$, where $|\Sigma_k| = 1$ for each $k = 1, \ldots, K$. For estimation of the covariance

parameters $\Sigma \in \mathcal{S}_{\mathbf{p}}^+$, we consider the following generalization of Stein's loss, which we call "multiway Stein's loss":

$$L_M(\Sigma, S) = \frac{s^2}{\sigma^2} \sum_{k=1}^{K} \frac{p}{p_k} \mathrm{tr}\left[S_k \Sigma_k^{-1}\right] - Kp \log\left(\frac{s^2}{\sigma^2}\right) - Kp, \quad \Sigma, S \in \mathcal{S}_{\mathbf{p}}^+. \tag{11}$$

It is easy to see that for $K = 1$, multiway Stein's loss reduces to Stein's loss. Multiway Stein's loss also has the attractive property of being invariant under multilinear transformations. To see this, define $SL_{\mathbf{p}}$ to be the set of lists of the form $A = (a, A_1, \ldots, A_K)$ for which $a > 0$ and $A_k \in SL_{p_k}$ for each k, with SL_{p_k} being the special linear group of $p_k \times p_k$ matrices with unit determinant. For two elements A and B of $SL_{\mathbf{p}}$, define $AB = (ab, A_1 B_1, \ldots, A_K B_K)$ and $A^T = (a, A_1^T, \ldots, A_K^T)$. Multiway Stein's loss is invariant under transformations of the form $\Sigma \to A\Sigma A^T$, as

$$\begin{aligned}
&L_M\left(A\Sigma A^T, ASA^T\right) \\
&= \frac{a^2 s^2}{a^2 \sigma^2} \sum_{k=1}^{K} \frac{p}{p_k} \mathrm{tr}\left[A_k S_k A_k^T \left(A_k \Sigma_k A_k^T\right)^{-1}\right] - Kp \log\left(\frac{a^2 s^2}{a^2 \sigma^2}\right) - Kp \\
&= \frac{s^2}{\sigma^2} \sum_{k=1}^{K} \frac{p}{p_k} \mathrm{tr}\left[S_k \Sigma_k^{-1}\right] - Kp \log\left(\frac{s^2}{\sigma^2}\right) - Kp \\
&= L_M(\Sigma, S).
\end{aligned}$$

In particular, (11) is invariant under $\mathcal{G}_{\mathbf{p}}^+$, as $\mathcal{G}_{\mathbf{p}}^+ \subset SL_{\mathbf{p}}$. Therefore, the best $\mathcal{G}_{\mathbf{p}}^+$-equivariant estimator under multiway Stein's loss can be obtained using Corollary 1.

Proposition 1. *(UMREE under multiway Stein's loss) Let*

$$\mathcal{E}_k = \left(E\left[\left(\sigma^2 \Sigma_k\right)^{-1} \middle| X\right]\right)^{-1},$$

where the expectation is with respect to the posterior distribution given by Equation 7. The minimizer of the posterior expectation

$$E\left[\frac{s^2}{\sigma^2} \sum_{k=1}^{K} \frac{p}{p_k} \mathrm{tr}\left[S_k^T \Sigma_k^{-1}\right] - Kp \log\left(\frac{s^2}{\sigma^2}\right) - Kp \middle| X\right]$$

with respect to s and the S_k's is

$$\begin{aligned}
\hat{\Sigma}_k &= \mathcal{E}_k / |\mathcal{E}_k|^{1/(p_k)} \\
\hat{\sigma}^2 &= \left(\sum_{k=1}^{K} \frac{1}{K} |\mathcal{E}_k|^{-1/p_k}\right)^{-1}.
\end{aligned}$$

The posterior expectation $E[(\sigma^2\Sigma_k)^{-1}|X]$ may be approximated by the Gibbs sampler of Section 3. That is, if $(\sigma^2\Sigma_k)^{(1)}, \ldots, (\sigma^2\Sigma_k)^{(T)}$ is a long sequence of values of $(\sigma^2\Sigma_k)$ simulated from the Gibbs sampler, then

$$\mathrm{E}[(\sigma^2\Sigma_k)^{-1}|X] \approx \sum_{t=1}^{T}[(\sigma^2\Sigma_k)^{(t)}]^{-1}/T.$$

The form of multiway Stein's loss (11) includes a weighted sum of $\mathrm{tr}(S_k\Sigma_k^{-1})$, $k = 1, \ldots, K$. We note that equivariant estimation of Σ is largely unaffected by changes to the weights in this sum:

Proposition 2. *Define weighted multiway Stein's loss as*

$$L_W(\Sigma, S) = \frac{s^2}{\sigma^2}\sum_{k=1}^{K}\frac{w_k}{p_k}\mathrm{tr}\left[S_k\Sigma_k^{-1}\right] - \left(\sum_{k=1}^{K}w_k\right)\log\left(\frac{s^2}{\sigma^2}\right) - \sum_{k=1}^{K}w_k,$$

for known $w_k > 0$, $k = 1, \ldots, K$. Then the UMREE under L_W is given by

$$\begin{aligned}\hat{\Sigma}_k &= \mathcal{E}_k/|\mathcal{E}_k|^{1/(p_k)} \\ \hat{\sigma}^2 &= \left(\sum_{k=1}^{K}\frac{w_k}{\sum_{i=1}^{K}w_i}|\mathcal{E}_k|^{-1/p_k}\right)^{-1}.\end{aligned}$$

The proof is very similar to that of Proposition 1 and is omitted. This proposition states that only estimation of the scale is affected when we "weight" the loss more heavily for some components of Σ than others.

The posterior distribution may also be used to obtain the UMREE under Stein's original loss L_S, as it too is invariant under transformations of the lower triangular product group. However, risk minimization with respect to L_S requires additional numerical approximations: Let $\mathcal{K}$ be the unique symmetric square root of $E[(\Sigma_K^{-1}\otimes\cdots\otimes\Sigma_1^{-1})/\sigma^2|X]$. This $\mathcal{K}$ may be approximated by the Gibbs sampler described in Section 3. Minimization of the risk with respect to L_S is equivalent to the minimization in $(s^2, S_1, \ldots, S_K)$ of

$$\begin{aligned}E[L_S(S,\Sigma)|X] &= s^2\mathrm{tr}\left(\mathcal{K}\left(S_K\otimes\cdots\otimes S_1\right)\mathcal{K}\right) - p\log\left(s^2\right) + c(\Sigma) \\ &= s^2\left|\left|\tilde{\mathcal{K}}\times\left\{S_1^{1/2}, \ldots, S_K^{1/2}, I_p\right\}\right|\right|^2 - p\log\left(s^2\right) + c(\Sigma) \\ &= \mathrm{tr}\left(s^2S_k\tilde{\mathcal{K}}_{(k)}S_{-k}\tilde{\mathcal{K}}_{(k)}^T\right) - p\log\left(|s^2S_k|\right)/p_k + c(\Sigma),\end{aligned}$$

where $\tilde{\mathcal{K}} \in \mathbb{R}^{p_1\times\cdots\times p_K\times p}$ is the array such that $\tilde{\mathcal{K}}_{(K+1)} = \mathcal{K}$, and $S_k^{1/2}$ is any square root matrix of S_k. Iteratively setting $s^2S_k = (\tilde{\mathcal{K}}_{(k)}S_{-k}\tilde{\mathcal{K}}_{(k)}^T)^{-1}p/p_k$ will decrease the posterior expected loss at

each step. This procedure is analogous to using the iterative flip-flop algorithm to find the MLE based on a sample covariance matrix of $E[(\Sigma_K^{-1} \otimes \cdots \otimes \Sigma_1^{-1})/\sigma^2 | X]$. Application of the results from [Wiesel, 2012b] show that the posterior risk has a property known as geodesic convexity, implying that any local minimizer obtained from this algorithm will also be a global minimizer.

4.2 An orthogonally equivariant estimator

The estimator in Proposition 1 depends on the ordering of the indices, and so it is not permutation equivariant. Mirroring the ideas studied in Takemura [1983], in this section we derive a minimax orthogonally equivariant estimator (which is necessarily permutation equivariant) that dominates the UMREE of Proposition 1. First, notice that by transforming the data and then back-transforming the estimator, we can obtain an estimator whose risk is equal to that of the UMREE: For $\Gamma = (1, \Gamma_1, \ldots, \Gamma_K) \in \{1\} \times \mathcal{O}_{p_1} \times \cdots \times \mathcal{O}_{p_K}$, where $\mathcal{O}_{p_k}$ is the group of p_k by p_k orthogonal matrices, let $\tilde{X} = X \times \{\Gamma_1, \ldots, \Gamma_K\}$. Then $\hat{\Sigma}(\tilde{X})$ is an estimator of $\Gamma\Sigma\Gamma^T$ and $\tilde{\Sigma}(X) = \Gamma^T \hat{\Sigma}(\tilde{X})\Gamma$ is an estimator of Σ. The risk of this estimator is the same as that of the UMREE $\hat{\Sigma}(X)$:

$$\begin{aligned}
R\left(\Sigma, \tilde{\Sigma}(X)\right) &= E\left[L_M\left(\Sigma, \Gamma^T \hat{\Sigma}(\tilde{X})\Gamma\right) \middle| \Sigma \right] \\
&= E\left[L_M\left(\Gamma\Sigma\Gamma^T, \hat{\Sigma}(\tilde{X})\right) \middle| \Sigma \right] \\
&= E\left[L_M\left(\Gamma\Sigma\Gamma^T, \hat{\Sigma}(X)\right) \middle| \Gamma\Sigma\Gamma^T \right] \\
&= R\left(\Gamma\Sigma\Gamma^T, \hat{\Sigma}(X)\right) \\
&= R\left(\Sigma, \hat{\Sigma}(X)\right)
\end{aligned}$$

where the second equality follows from the invariance of the loss, the third equality follows from a change of variables, and the last equality follows because the risk of $\hat{\Sigma}$ is constant over the parameter space. The UMREE $\hat{\Sigma}$ and the estimator $\tilde{\Sigma}$ have the same risks but are different. Since multiway Stein's loss is convex in each argument, averaging these estimators somehow should produce a new estimator that dominates them both.

In the multivariate normal case in which $K = 1$, averaging the value of $\Gamma^T \hat{\Sigma}(\Gamma X)\Gamma$ with respect to the uniform (invariant) measure for Γ over the orthogonal group results in the estimator of Takemura [1983]. This estimator is orthogonally equivariant, dominates the UMREE and is therefore also minimax. Constructing an analogous estimator in the multiway case is more complicated, as

it is not immediately clear how the back-transformed estimators should be averaged. Direct numerical averaging of estimates of $\sigma^2(\Sigma_1 \otimes \cdots \otimes \Sigma_K)$ will generally produce an estimate that is not separable and therefore outside of the parameter space. Similarly, averaging estimates of each Σ_k separately will not work, as the space of covariance matrices with determinant one is not convex.

Our solution to this problem is to average a transformed version of $\Sigma = (\sigma^2, \Sigma_1, \ldots, \Sigma_K)$ for which each Σ_k lies in the convex set of trace-1 covariance matrices, then transform back to our original parameter space. The resulting estimator, which we call the multiway Takemura estimator (MWTE), is orthogonally equivariant and uniformly dominates the UMREE.

Proposition 3. *Let $\hat{\sigma}^2(\Gamma, X)$ and $\hat{\Sigma}_k(\Gamma, X)$ be the UMREEs of σ^2 and $\Gamma_k \Sigma_k \Gamma_k^T$ based on data $X \times \{\Gamma_1, \ldots, \Gamma_K, I_n\}$. Let*

$$S_k(X) = \int_{\mathcal{O}_{p_K}} \cdots \int_{\mathcal{O}_{p_1}} \frac{\Gamma_k^T \hat{\Sigma}_k(\Gamma, X) \Gamma_k}{\operatorname{tr}\left(\hat{\Sigma}_k(\Gamma, X)\right)} \, d\Gamma_1 \cdots d\Gamma_K$$

and

$$\tilde{\sigma}^2(X) = \int_{\mathcal{O}_{p_K}} \cdots \int_{\mathcal{O}_{p_1}} \hat{\sigma}^2(\Gamma, X) \; d\Gamma_1 \cdots d\Gamma_K.$$

Let $\tilde{\Sigma}_k(X) = S_k(X)/|S_k(X)|^{1/p_k}$ for $k = 1, \ldots, K$. Then $(\tilde{\sigma}^2(X), \tilde{\Sigma}_1(X), \ldots, \tilde{\Sigma}_K(X))$ is orthogonally equivariant and uniformly dominates the UMREE of Proposition 1.

Note that "averaging" over any subset of $\mathcal{O}_{p_1} \times \cdots \times \mathcal{O}_{p_K}$ in the manner of Proposition 3 will uniformly decrease the risk. By averaging with respect to the uniform measure over the orthogonal group, we obtain an estimator that has the attractive property of being orthogonally equivariant.

In practice it is computationally infeasible to integrate over the space of orthogonal matrices. However, we may obtain a stochastic approximation to the MWTE as follows: Independently for each $t = 1, \ldots, T$ and $k = 1, \ldots, K$, simulate $\Gamma_k^{(t)}$ from the uniform distribution on $\mathcal{O}_{p_k}$. Let

$$S_k(X) = \frac{1}{T} \sum_{t=1}^{T} \frac{{\Gamma_k^{(t)}}^T \hat{\Sigma}_k\left(\Gamma^{(t)}, X\right) \Gamma_k^{(t)}}{\operatorname{tr}\left(\hat{\Sigma}_k\left(\Gamma^{(t)}, X\right)\right)} \;,\; \tilde{\sigma}_T(X) = \frac{1}{T} \sum_{t=1}^{T} \hat{\sigma}(\Gamma^{(t)}, X).$$

Set $\tilde{\Sigma}_{k,T}(X) = S_k(X)/|S_k(X)|^{1/p_k}$ for $k = 1, \ldots, K$. Then an approximation to the MWTE is

$$\tilde{\Sigma}_T = \left(\tilde{\sigma}_T^2(X), \tilde{\Sigma}_{1,T}(X), \ldots, \tilde{\Sigma}_{K,T}(X)\right). \tag{12}$$

This is a randomized estimator which is orthogonally invariant in the sense of Definition 6.3 of Eaton [1989].

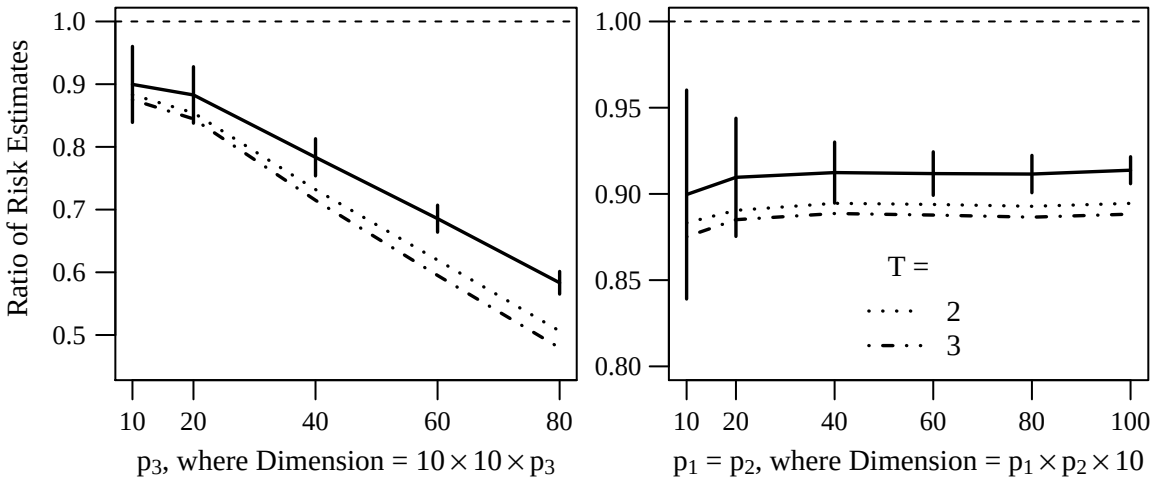

Figure 1: Risk comparisons for the MLE, UMREE and MWTE. Both panels plot Monte Carlo estimates of the risk ratios of the UMREE to the MLE in solid lines, and the approximate MWTE to the MLE in dashed lines. The width of the vertical bars is one standard deviation of the ratio of the UMREE loss to the MLE loss, across the 100 data sets.

4.3 Simulation results

We numerically compared the risks of the MLE, UMREE, and the MWTE under several three-way array normal distributions, using a variety of values of (p_1, p_2, p_3) and with $n = 1$. For each (p_1, p_2, p_3) under consideration, we simulated 100 data arrays from the array normal model. As the risk of both the MLE and the UMREE are constant over the parameter space, it is sufficient to compare their risks at a single point in the parameter space, which we took to be $\Sigma = (1, I_{p_1}, I_{p_2}, I_{p_3})$. Risks were approximated by averaging the losses of each estimator across the 100 simulated data arrays. For each data array, the MLE was obtained from the iterative coordinate descent algorithm outlined in [Hoff, 2011]. Each UMREE was approximated based on 1250 iterations of the Gibbs sampler described in Section 3, from which the first 250 iterations were discarded to allow for convergence to the stationary distribution (convergence appeared to be essentially immediate).

The ratio of risk estimates across several values of (p_1, p_2, p_3) are are plotted in solid lines in Figure 1. We considered array dimensions in which the first two dimensions were identical. This scenario could correspond to, for example, data arrays representing longitudinal relational or network measurements between $p_1 = p_2$ nodes at p_3 time points. The first panel of the figure considers

the relative performance of the estimators as the "number of time points" (p_3) increases. The results indicate that the UMREE provides substantial and increasing risk improvements compared to the MLE as p_3 increases. However, the right panel indicates that the gains are not as dramatic and not increasing when the "number of nodes" ($p_1 = p_2$) increases while p_3 remains fixed. Even so, the variability in the ratio of losses (shown with vertical bars) decreases as the number of nodes increases, indicating an increasing probability that that the UMREE will beat the MLE in terms of loss.

We also compared these risks to the risk of the approximate MWTE given in (12), with $T \in \{2, 3\}$. The risks for the approximate MWTE relative to those of the MLE are shown in dashed lines in the two panels of the Figure, and indicate non-trivial improvements in risk as compared to the UMREE. We examined values of T greater than 3 but found no appreciable further reduction in the risk. Note, however, that the MWTE does not have constant risk over the parameter space (though MWTE will have constant risk over the orbits of the orthogonal product group).

5 Discussion

This article has extended the results of James and Stein [1961] and Takemura [1983] by developing equivariant and minimax estimators of the covariance parameters in the array normal model. Considering the class of estimators equivariant with respect to a special lower triangular group, we showed that the uniform minimum risk equivariant estimator (UMREE) can be viewed as a generalized Bayes estimator that can be obtained from a simple Gibbs sampler. We obtained an orthogonally equivariant estimator based on this UMREE by combining values of the UMREE under orthogonal transformations of the data. Both the UMREE and the orthogonally equivariant estimator are minimax, and both dominate any unique MLE in terms of risk.

Empirical results in Section 4 indicate that the risk improvements of the UMREE over the MLE can be substantial, while the improvements of the orthogonally equivariant estimator over the UMREE are more modest. However, the risk improvements depend on the array dimensions in a way that is not currently understood. Furthermore, we do not yet know the minimal conditions necessary for the propriety of the posterior or the existence of the UMREE. Empirical results from the simulations in Section 4 suggest that the UMREE exists for sample sizes as low as $n = 1$, at least for the array dimensions in the study. This is similar to the current state of knowledge for the

existence of the MLE: The array normal likelihood is trivially bounded for $n \geq p$ (as it is bounded by the maximized likelihood under the unconstrained p-variate normal model), and some sufficient conditions for uniqueness of the MLE are given in Ohlson et al. [2013]. However, empirical results (not shown) suggest that a unique MLE may exist for $n = 1$ for some array dimensions (although not for others). Obtaining necessary and sufficient conditions for the existence of the UMREE and the MLE is an ongoing area of research of the authors.

A Proofs

A.1 Proof of Theorem 1

Proof. Let $a > 0$, $A_k \in \mathcal{G}_{p_k}^+$ for all $k = 1, \ldots, K$. Let t be a fixed element in $\mathbb{R}^+$ and T_k be fixed elements in $\mathcal{G}_{p_k}^+$ for $k = 1, \ldots, K$. In the terminology of Definition 1.7 of Eaton [1989], the integral with respect to Lebesgue measure is relatively right invariant with multiplier

$$\chi\left(t, T_1, \ldots, T_K\right) = t \prod_{k=1}^{K} \prod_{i=2}^{p_k} T_{k[i,i]}^{2-i}$$

if the following holds:

$$\begin{aligned} &\int_{\mathcal{G}_{\mathbf{p}}^+} f(a/t, A_1 T_1^{-1}, \ldots, A_K T_K^{-1}) d\mu\left(a, A_1, \ldots, A_K\right) \\ &= \left(t \prod_{k=1}^{K} \prod_{i=2}^{p_k} T_{k[i,i]}^{2-i}\right) \int_{\mathcal{G}_{\mathbf{p}}^+} f(a, A_1, \ldots, A_K) d\mu\left(a, A_1, \ldots, A_K\right), \end{aligned} \tag{13}$$

for arbitrary $f()$. If (13) holds, then by Theorem 1.6 of Eaton [1989], a right invariant measure over the group $\mathcal{G}_{\mathbf{p}}^+$ is

$$\chi\left(a, A_1, \ldots, A_K\right)^{-1} = \frac{1}{a} \prod_{k=1}^{K} \prod_{i=2}^{p_k} A_{k[i,i]}^{i-2} d\mu.$$

It remains to make a change of variables to show that (13) holds. For $E_k, T_k \in \mathcal{G}_{p_k}^+$ with T_k fixed for $k = 1, \ldots, K$, let $g_k(E_k) = E_k T_k$ for $k = 1, \ldots, K$. For $e, t > 0$ with t fixed let $g(e) = et$. The Jacobian for transforming the scale, $g(e) = et$, is t. The Jacobian for the transformation $g_k(E_k) = E_k T_k$ is

$$J(E_k) = \prod_{i=2}^{p_k} T_{k[i,i]}^{2-i}. \tag{14}$$

To see this, note that this transformation is equivalent to $p_k(p_k+1)/2-1$ linear transformations of the form:

$$g_{i,j} : E_{k[i,j]} \mapsto \sum_{j \le m \le i} E_{k[i,m]} T_{k[m,j]} \text{ for all } 1 \le j \le i \le p_k \text{ s.t. } (i,j) \neq (1,1).$$

Stack the elements of E_k into the following vector:

$$\begin{aligned} s = &(E_{k[p_k,p_k]}, E_{k[p_k,p_k-1]}, E_{k[p_k-1,p_k-1]}, E_{k[p_k,p_k-2]}, \\ &E_{k[p_k-1,p_k-2]}, E_{k[p_k-2,p_k-2]}, E_{k[p_k,p_k-3]}, \ldots, E_{k[2,1]}), \end{aligned}$$

and notice that the matrix of the linear transformation is lower triangular:

$$u = \begin{pmatrix} T_{k[p_k,p_k]} & 0 & 0 & 0 & \cdots & \\ T_{k[p_k,p_k-1]} & T_{k[p_k-1,p_k-1]} & 0 & 0 & & \\ 0 & 0 & T_{k[p_k-1,p_k-1]} & 0 & & \\ T_{k[p_k,p_k-2]} & T_{k[p_k-1,p_k-2]} & 0 & T_{k[p_k-2,p_k-2]} & & \\ \vdots & & & & \ddots & \\ \vdots & & & & & \ddots \\ 0 & \cdots & \cdots & \cdots & T_{k[2,1]} & \cdots T_{k[1,1]} \end{pmatrix}$$

where in the diagonal, each $T_{k[i,i]}$ is repeated $p_k - i + 1$ times for $i = 2, 3, \ldots, p_k$, and $T_{k[1,1]}$ is repeated $p_k - 1$ times. That is, the linear transformation can be written as:

$$g_k(s) = us.$$

Hence the determinant of the Jacobian is

$$|u| = T_{k[1,1]}^{p_k-1} \prod_{i=2}^{p_k} T_{k[i,i]}^{p_k-i+1} = \prod_{i=2}^{p_k} T_{k[i,i]}^{2-i},$$

where the second equality results from our parameterization of $\mathcal{G}_{p_k}^+$,

$$\prod_{i=2}^{p_k} T_{k[i,i]}^{-1} = T_{k[1,1]}.$$

□

A.2 Proof of Theorem 2

Consider the reformulation of the problem to a parameterization of $\Sigma = \sigma^2(\Psi_K \Psi_K^T \otimes \cdots \otimes \Psi_1 \Psi_1^T)$ where $\Psi_{k[1,1]} = 1$ for $k = 1, \ldots, K$. That is, we now work with the group $\mathcal{G}^1_{\mathbf{p}} = \{(a, A_1, \ldots, A_K) | a > 0, A_k \in \mathcal{G}^1_{p_k} \text{ for } k = 1, \ldots, K\}$ where $\mathcal{G}^1_{p_k}$ is the group of p_k by p_k lower triangular matrices with positive diagonal elements and 1 in the $(1,1)$ position. The group operation in $\mathcal{G}^1_{p_k}$ is matrix multiplication, and that of $\mathcal{G}^1_{\mathbf{p}}$ is component-wise multiplication. The left and right Haar measures over $\mathcal{G}^1_{p_k}$ are easy to derive:

Lemma 5. *For $E_k, T_k \in \mathcal{G}^1_{p_k}$ with T_k fixed, the Jacobian for the transformation $g(E_k) = E_k T_k$ is*

$$J(E_k) = \prod_{i=2}^{p_k} T_{k[i,i]}^{p_k - i + 1}$$

the Jacobian for the transformation $g(E_k) = T_k E_k$ is

$$J(E_k) = \prod_{i=2}^{p_k} T_{k[i,i]}^{i}$$

So the right Haar measure is

$$d\nu_r(E_k) = \prod_{i=2}^{p_k} E_{k[i,i]}^{-p_k + i - 1}$$

Proof. The proof is very similar to those in Propositions 5.13 and 5.14 of Eaton [1983], noting that $T_{k[1,1]} = 1$. □

We'll eventually need the inverse transformation, which follows directly from Theorem 3 of chapter 8 section 4 of Magnus and Neudecker [1988].

Lemma 6. *For $E_k \in \mathcal{G}^1_{p_k}$, the Jacobian for the transformation $g(E_k) = E_k^{-1}$ is*

$$\prod_{i=2}^{p_k} E_{k[i,i]}^{-p_k - 1} \tag{15}$$

Proof. From Magnus and Neudecker [1988], $d(E_k^{-1}) = -E_k^{-1}(dE_k)E_k^{-1}$. Using Lemma 5, the Jacobian of the first transformation, $g_1(dE_k) = E_k^{-1}(dE_k)$ is $\prod_{i=2}^{p_k} E_{k[i,i]}^{-i}$. Jacobian of the second transformation $g_2(dE_k) = (dE_k)E_k^{-1}$ is $\prod_{i=2}^{p_k} E_{k[i,i]}^{-p_k + i - 1}$. Hence, overall Jacobian is (15). □

Under this new parameterization, the likelihood is

$$
\begin{aligned}
&p(X|\sigma, \Psi_1, \ldots, \Psi_K) \\
&= (2\pi)^{np/2}|\sigma^2(\Psi_K\Psi_K^T \otimes \cdots \otimes \Psi_1\Psi_1^T)|^{-n/2} \\
&\times \exp\{-||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^T/(2\sigma^2)\}\} \\
&\propto \sigma^{-np} \prod_{k=1}^{K}\prod_{i=2}^{p_k} \Psi_{k[i,i]}^{-np/p_k} \exp\{-||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^T/(2\sigma^2)\}\},
\end{aligned}
$$

where $p = \prod_{k=1}^{K} p_k$. The (improper) prior is

$$
\pi(\sigma, \Psi_1, \ldots, \Psi_K) \propto \frac{1}{\sigma} \prod_{k=1}^{K}\prod_{i=2}^{p_k} \Psi_{k[i,i]}^{i-p_k-1}.
$$

Hence, the posterior is

$$
\sigma^{-np-1} \prod_{k=1}^{K}\prod_{i=2}^{p_k} \Psi_{k[i,i]}^{i-np/p_k-p_k-1} \exp\{-||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^T/(2\sigma^2)\}.
$$

Since $\sigma^2|\Psi \sim$ inverse-gamma$(np/2, ||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^2/2)$, we can integrate out σ^2, obtaining

$$
\pi(\Psi_1, \ldots, \Psi_K|X) \propto ||X \times \{\Psi_1^{-1}, \ldots, \Psi_K^{-1}, I_n\}||^{-np} \prod_{k=1}^{K}\prod_{i=2}^{p_k} \Psi_{k[i,i]}^{i-np/p_k-p_k-1}.
$$

Let $S = X_{(K+1)}^T X_{(K+1)}$, the sample covariance matrix, then

$$
\begin{aligned}
&\pi(\Psi_1, \ldots, \Psi_K|X) \\
&\propto \mathrm{tr}[S(\Psi_K^{-T}\Psi_K^{-1} \otimes \cdots \otimes \Psi_1^{-T}\Psi_1^{-1})]^{-np/2} \prod_{k=1}^{K}\prod_{i=2}^{p_k} \Psi_{k[i,i]}^{i-np/p_k-p_k-1}.
\end{aligned}
$$

Let $L_k = \Psi_k^{-1}$ for $k = 1, \ldots, K$. Then, using Lemma 6, we have

$$
\begin{aligned}
&\pi(L_1, \ldots, L_K|X) \\
&\propto \mathrm{tr}[S(L_K^T L_K \otimes \cdots \otimes L_1^T L_1)]^{-np/2} \prod_{k=1}^{K}\prod_{i=2}^{p_k} L_{k[i,i]}^{np/p_k-i}
\end{aligned} \tag{16}
$$

The posterior density is integrable if and only if (16) is integrable. We will now prove that when $n > \prod_{k=1}^{K} p_k$ then (16) is integrable. First, consider, consider the integral over $\mathcal{G}_p^1$, where $p = \prod_{k=1}^{K} p_k$,

$$
\int_{\mathcal{G}_p^1} \mathrm{tr}\left(VSV^T\right)^{-np/2} \prod_{i=2}^{p} V_{[i,i]}^{np-p-1} dV \tag{17}
$$

Let $e = (1, 0, \ldots, 0)^T$, the vector of length p with a 1 in the first position and 0's everywhere else. Then $V = (e^T, V_2^T)^T$ and

$$\begin{aligned}\text{tr}(VSV^T) &= tr(e_1^T S e_1) + tr(V_2 S V_2^T) = S_{[1,1]} + tr(V_2 S V_2^T) \\ &= (1 + tr(V_2 S V_2^T)/S_{[1,1]}) S_{[1,1]} = (1 + tr(V_2 S_T (V_2 S_T)^T)/S_{[1,1]}) S_{[1,1]},\end{aligned}$$

where $S = S_T S_T^T$ is the lower triangular Cholesky decomposition of S. Let $W = V_2 S_T$, so $V_2 = W S_T^{-1}$. The Jacobian of this transformation is $S_{T[1,1]}^{1-p} \prod_{i=2}^{p} S_{T[i,i]}^{i-p-1}$ (same as the Jacobian in Proposition 5.14 of Eaton [1983] except we have one less $S_{T[1,1]}$ term). Then Equation (17) is proportional to

$$\begin{aligned}&\int_{\mathcal{G}_p^1} \left((1 + \text{tr}(WW^T)/S_{[1,1]}\right)^{-np/2} \prod_{i=2}^{p} W_{[i,i]}^{np-p-1} dW \\ &= \int_{\mathcal{G}_p^1} \left((1 + \mathbf{w} D \mathbf{w}/(np-p)\right)^{-(np-p+p)/2} \prod_{i=2}^{p} W_{[i,i]}^{np-p-1} dW,\end{aligned}$$

where $\mathbf{w}$ is a vector containing all the non-zero elements of W and $D = (n-p) I_p / S_{[1,1]}$. Notice that $\left((1 + \mathbf{w} D \mathbf{w}/(np-p)\right)^{-(np-p+p)/2}$ is the kernel of a multivariate T distribution with degrees of freedom $np - p$ and scale matrix $D^{-1} = S_{[1,1]} I_p/(np-p)$ [Kotz and Nadarajah, 2004, equation (1.1)]. Note that $E[W_{[i,j]}^{\nu}] < \infty$ if $\nu < n - p$ [Kotz and Nadarajah, 2004, section 1.7]. In particular, $n - p - 1 < n - p$. Hence

$$\int_{\mathcal{G}_p^1} \text{tr}\left(VSV^T\right)^{-np/2} \prod_{i=2}^{p} V_{[i,i]}^{np-p-1} dV < \infty$$

Using this, we have the following inequalities:

$$\begin{aligned}\infty &> \int_{\mathcal{G}_p^1} \text{tr}\left[VSV^T\right]^{-np/2} \prod_{i=1}^{p} V_{k[i,i]}^{np-p-1} dV \\ &= \int_{\mathcal{G}_p^1} \text{tr}\left[VSV^T\right]^{-np/2} |V|^{np-p-1} dV \\ &\geq \int_{\mathcal{G}_{p_1}^1 \times \cdots \times \mathcal{G}_{p_K}^1} \text{tr}\left[(L_K \otimes \cdots \otimes L_1) S (L_K \otimes \cdots \otimes L_1)^T\right]^{-np/2} \\ &\times |L_K \otimes \cdots \otimes L_1|^{np-p-1} dL_1 \cdots dL_K \\ &= \int_{\mathcal{G}_{p_1}^1 \times \cdots \times \mathcal{G}_{p_K}^1} \text{tr}\left[S(L_K^T L_K \otimes \cdots \otimes L_1^T L_1)\right]^{-np/2} \\ &\times \prod_{k=1}^{K} \prod_{i=2}^{p_k} L_{k[i,i]}^{(np-p-1)p/p_k} dL_1 \cdots dL_K,\end{aligned}$$

where the second inequality results from integrating over a smaller space. Note the following results: (1) $(np-p-1)p/p_k \geq np/p_k - i_k$ for all $k = 1, \ldots, K$ and $i_k = 2, \ldots, p_k$ if $n \geq p$, (2) $L_{k[i,i]} > 0$, and (3) $E[|X|^{r_1}] < \infty$ and $r_1 > r_2 \Rightarrow E[|X|^{r_2}] < \infty$. Hence,

$$\begin{aligned}\infty > &\int_{\mathcal{G}^1_{p_1} \times \cdots \times \mathcal{G}^1_{p_K}} \operatorname{tr}[S(L_K^T L_K \otimes \cdots \otimes L_1^T L_1)]^{-np/2} \\ &\times \prod_{k=1}^{K} \prod_{i=2}^{p_k} L_{k[i,i]}^{np/p_k - i} dL_1 \cdots dL_K\end{aligned}$$

and the result is proved.

A.3 Proof of Lemma 1

Proof. Let VV^T be the lower triangular Cholesky decomposition of a Wishart$_p(\nu, I_p)$-distributed random matrix. Recall from Bartlett's decomposition [Bartlett, 1933] that the elements of V are independent with

$$V_{[i,i]}^2 \sim \chi_{\nu-i+1}^2 \text{ and } V_{[i,j]} \sim N(0,1).$$

Let $S = V^T V$. For $i \neq j$, we have

$$E\left[S_{[i,j]}\right] = E\left[\sum_{k=1}^{p} V_{[k,i]} V_{[k,j]}\right] = \sum_{k=1}^{p} E\left[V_{[k,i]}\right] E\left[V_{[k,j]}\right].$$

For $i \neq j$, we have either $E[V_{[k,i]}] = 0$ or $E[V_{[k,j]}] = 0$ for all $k = 1, \ldots, p$. Hence, $E[S_{[i,j]}] = 0$ for all $i \neq j$.

For $i = j$, we have

$$\begin{aligned}E\left[S_{[i,j]}\right] &= E\left[\sum_{k=1}^{p} V_{[k,i]} V_{[k,j]}\right] = \sum_{k=1}^{p} E\left[V_{[k,i]}^2\right] = E\left[V_{[i,i]}^2\right] + \sum_{k=i+1}^{p} E\left[V_{[k,i]}^2\right] \\ &= \nu - i + 1 + \sum_{k=i+1}^{p} 1 = \nu - i + 1 + p - i = \nu + p + 1 - 2i.\end{aligned}$$

This expectation has been calculated in other papers [James and Stein, 1961, Eaton et al., 1987, for example]. □

A.4 Proof of Lemma 2

Proof. We proceed by invariance arguments. The Jacobian, $J(\sigma, \Psi)$, is the unique continuous function that satisfies

$$\begin{aligned}\int_{G_{p_k}^+} f(L)\frac{dL}{\prod_{i=1}^{p_k} L_{[i,i]}^{p_k-i+1}} &= \int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(\sigma\Psi)\frac{J(\sigma,\Psi)d\sigma d\Psi}{\prod_{i=1}^{p_k}(\sigma\Psi_{[i,i]})^{p_k-i+1}} \\ &= \int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(\sigma\Psi)\frac{J(\sigma,\Psi)d\sigma d\Psi}{\sigma^{p_k(p_k+1)/2}\prod_{i=1}^{p_k}\Psi_{[i,i]}^{p_k-i+1}},\end{aligned}$$

where $dL/(\prod_{i=1}^{p_k} L_{[i,i]}^{p_k-i+1})$ is a right invariant measure with respect to the action $L \mapsto LA$ on $G_{p_k}^+$ for $A \in G_{p_k}^+$ [Eaton, 1983, Proposition 5.14]. Hence, this invariance property must also hold for the right integral. So for $b > 0$ and $B \in \mathcal{G}_{p_k}^+$, we have that $bB \in G_{p_k}^+$ and

$$\begin{aligned}&\int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(\sigma\Psi)\frac{J(\sigma,\Psi)d\sigma d\Psi}{\sigma^{p_k(p_k+1)/2}\prod_{i=1}^{p_k}\Psi_{[i,i]}^{p_k-i+1}} \\ &= \int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(b\sigma\Psi B)\frac{J(\sigma,\Psi)d\sigma d\Psi}{\sigma^{p_k(p_k+1)/2}\prod_{i=1}^{p_k}\Psi_{[i,i]}^{p_k-i+1}}.\end{aligned}$$

So making the change of variables $\sigma = e/b$ and $\Psi = EB^{-1}$, we have

$$\begin{aligned}&\int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(b\sigma\Psi B)\frac{J(\sigma,\Psi)d\sigma d\Psi}{\sigma^{p_k(p_k+1)/2}\prod_{i=1}^{p_k}\Psi_{[i,i]}^{p_k-i+1}} \\ &= \int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(eE)\frac{\frac{1}{b}\prod_{i=2}^{p_k} B_{[i,i]}^{i-2} J(e/b, EB^{-1})dedE}{(e/b)^{p_k(p_k+1)/2}\prod_{i=1}^{p_k} E_{[i,i]}^{p_k-i+1} B_{[i,i]}^{i-p_k-1}} \\ &= \int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(eE)\frac{b^{p_k(p_k+1)/2-1} B_{[1,1]}^{p_k}\prod_{i=2}^{p_k} B_{[i,i]}^{p_k-1} J(e/b, EB^{-1})dedE}{e^{p_k(p_k+1)/2}\prod_{i=1}^{p_k} E_{[i,i]}^{p_k-i+1}} \\ &= \int_{\mathbb{R}\times\mathcal{G}_{p_k}^+} f(eE)\frac{b^{p_k(p_k+1)/2-1} B_{[1,1]} J(e/b, EB^{-1})dedE}{e^{p_k(p_k+1)/2}\prod_{i=1}^{p_k} E_{[i,i]}^{p_k-i+1}},\end{aligned}$$

where we used (14) for the first equality and our parameterization of $\mathcal{G}_{p_k}^+$, $\prod_{i=2}^{p_k} B_{[i,i]}^{-1} = B_{[1,1]}$, for the last equality. So we must have that

$$J(\sigma, \Psi) = b^{p_k(p_k+1)/2-1} B_{[1,1]} J(\sigma/b, \Psi B^{-1}).$$

Set $B = \Psi$ and $b = \sigma$ to obtain: $J(\sigma, \Psi) = \sigma^{p_k(p_k+1)/2-1}\Psi_{[1,1]} J(1, I)$, where $J(1, I)$ is a constant. □

A.5 Proof of Lemma 3

Let $S^{-1} \sim \text{Wishart}_p(\nu, I_p)$ and partition S^{-1} and $S \sim \text{inverse-Wishart}_p(\nu, I_p)$ conformably such that $p_1 + p_2 = p$:

$$S^{-1} = \begin{pmatrix} S^{11} & S^{12} \\ S^{21} & S^{22} \end{pmatrix}, \quad S = \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix}.$$

Denote $S^{11\bullet 2} = S^{11} - S^{12}(S^{22})^{-1}S^{21}$, the Schur complement. The following are well known properties of the Wishart distribution (see, for example, Proposition 8.7 of Eaton [1983])

$$\begin{aligned} S^{22} &\sim \text{Wishart}_{p_2}(I_{p_2}, \nu) \\ S^{21}|S^{22} &\sim N_{p_2 \times p_1}(0, S^{22} \otimes I_{p_1}) \\ S^{11\bullet 2} &\sim \text{Wishart}_{p_1}(I_{p_1}, \nu - p_2) \\ S^{11\bullet 2} &\text{ is independent of } \{S^{22}, S^{21}\} \end{aligned}$$

The relationship of the inverse of a partitioned matrix (see, for example, Section 0.7.3 of Horn and Johnson [2012]) implies that

$$S_{11} = (S^{11\bullet 2})^{-1} \sim \text{inverse-Wishart}_{p_1}(I_{p_1}, \nu - p_2) \tag{18}$$

$$S_{22\bullet 1} = (S^{22})^{-1} \sim \text{inverse-Wishart}_{p_2}(I_{p_2}, \nu) \tag{19}$$

$$\begin{aligned} S_{21}|S_{11}, S_{22\bullet 1} &\stackrel{d}{=} -(S^{22})^{-1}S^{21}(S^{11\bullet 2})^{-1} \\ &\sim N_{p_2 \times p_1}(0, (S^{22})^{-1} \otimes (S^{11\bullet 2})^{-1}(S^{11\bullet 2})^{-1}) \qquad (20) \\ &= N_{p_2 \times p_1}(0, S_{22\bullet 1} \otimes S_{11}S_{11}). \end{aligned}$$

It is also well known that

$$\text{if } p = 1 \text{ then } S \sim \text{ inverse-gamma}(\nu/2, 1/2). \tag{21}$$

We should be able to use these results to come up with the distribution of the elements of the lower triangular Cholesky decomposition from an inverse-Wishart distributed random matrix, which seems surprisingly difficult to find in the literature.

Proof of Lemma 3. We proceed by induction on the dimension. It is clearly true for $n = 1$. Assume it is true for $n - 1$. Then partition $S_{[1:n,1:n]} \sim \text{ inverse-Wishart}_n(I_n, \nu - p + n)$ such that the top

left submatrix, S_{11}, is $n-1$ by $n-1$.

$$S_{[1:n,1:n]} = \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & s_{22} \end{pmatrix} = \begin{pmatrix} W_1 & 0 \\ S_{21}W_1^{-T} & s_{22\bullet 1}^{1/2} \end{pmatrix} \begin{pmatrix} W_1^T & W_1^{-1}S_{12} \\ 0 & s_{22\bullet 1}^{1/2} \end{pmatrix}.$$

Note that $S_{11} = W_1 W_1^T$. Using (18)-(21), we have that:

$$\begin{aligned} W_{[n,n]}^2 = s_{22\bullet 1} &\sim \text{ inverse-gamma}\left((\nu - p + n)/2, 1/2\right) \\ S_{21}W_1^{-T}|W_1, s_{22\bullet 1} &= S_{21}S_{11}^{-1/2^T}|S_{11}, s_{22\bullet 1} \\ &\sim N_{1\times n-1}\left(0, \left(s_{22\bullet 1} \otimes W_1^T W_1\right)\right) \\ &= N_{n-1}(0, s_{22\bullet 1}W_1^T W_1) \\ &= N_{n-1}(0, W_{[n,n]}^2 W_1^T W_1). \end{aligned}$$

□

A.6 Proof of Lemma 4

Proof. We proceed by induction on the dimension. It is clearly true for $n = 1$. Assume it is true for $n-1$. Note that for lower triangular matrices, the $[1:n, 1:n]$ submatrix of the inverse is the inverse of the $[1:n, 1:n]$ submatrix. Hence, partition $W_{k[1:n,1:n]} = V_{k[1:n,1:n]}^{-1}$ by:

$$V_{k[1:n,1:n]} = \begin{pmatrix} V_{11} & 0 \\ V_{21} & v_{22} \end{pmatrix}, \quad W_{k[1:n,1:n]} = \begin{pmatrix} W_{11} & 0 \\ W_{21} & w_{22} \end{pmatrix},$$

where the top left submatrix is $n-1$ by $n-1$. Then $v_{22}^2 = 1/w_{22}^2$ is clearly $\chi_{\nu-n+1}^2$. We have that $V_{21} = -w_{22}^{-1}W_{21}W_{11\bullet 2}^{-1}$. Also, $W_{11\bullet 2} = W_{11} - W_{21} * 0/w_{22} = W_{11}$. Since

$$W_{21}|W_{11}, w_{22} \sim N_{n-1}\left(0, w_{22}^2 W_{11}^T W_{11}\right),$$

we have that

$$-w_{22}^{-1}W_{21}W_{11}^{-1}|W_{11}, w_{22} \sim N_{n-1}\left(0, I\right).$$

Hence, the result is proved. □

A.7 Proof of Proposition 1

Proof. This minimization problem is equivalent to minimizing

$$\begin{aligned} & s^2 \sum_{k=1}^{K} \frac{p}{p_k} \mathrm{tr}\left(S_k E\left[\left(\sigma^2 \Sigma_k\right)^{-1}\right]\right) - Kp \log\left(s^2\right) \\ & = s^2 \sum_{k=1}^{K} \frac{p}{p_k} \mathrm{tr}\left(S_k \mathcal{E}_k^{-1}\right) - Kp \log\left(s^2\right). \end{aligned}$$

Let us absorb the scale parameter into S_k. That is, let $\tilde{S}_k = s^2 S_k$, then $s^2 = |\tilde{S}_k|^{1/p_k}$, and we wish to minimize with respect to $\tilde{S}_k$:

$$\frac{p}{p_k} \mathrm{tr}\left(\tilde{S}_k \mathcal{E}_k^{-1}\right) - \frac{Kp}{p_k} \log\left(|\tilde{S}_k|\right) + |\tilde{S}_k|^{1/p_k} \sum_{j \neq k} \frac{p}{p_j} \mathrm{tr}\left(S_j^T \mathcal{E}_j^{-1}\right).$$

Letting $\lambda = \frac{p_k}{p} \sum_{j \neq k} \frac{p}{p_j} \mathrm{tr}\left(S_j^T \mathcal{E}_j^{-1}\right)$, this is equivalent to minimizing:

$$\mathrm{tr}\left(\tilde{S}_k \mathcal{E}_k^{-1}\right) - K \log\left(|\tilde{S}_k|\right) + |\tilde{S}_k|^{1/p_k} \lambda$$

with respect to $\tilde{S}_k$.

Since the mapping $\tilde{S}_k \mapsto \mathcal{E}_k^{-1/2} \tilde{S}_k \mathcal{E}_k^{-1/2} = \Omega$ is a bijection of the set of $p_k \times p_k$ symmetric positive definite matrices, we can write:

$$\begin{aligned} & \min_{\tilde{S}_k > 0} \left\{ \mathrm{tr}\left(\tilde{S}_k \mathcal{E}_k^{-1}\right) - K \log\left(|\tilde{S}_k|\right) + |\tilde{S}_k|^{1/p_k} \lambda \right\} \\ & = \min_{\Omega > 0} \left\{ \mathrm{tr}\left(\Omega\right) - K \log\left(|\Omega|\right) + |\Omega|^{1/p_k} \lambda^* + K \log(|\mathcal{E}_k|) \right\} \\ & = \min_{\omega_1 \geq \cdots \geq \omega_{p_k} > 0} \left\{ \sum_{i=1}^{p_k} \omega_i - K \sum_{i=1}^{p_k} \log(\omega_i) + \lambda^* \prod_{i=1}^{p_k} \omega_i^{1/p_k} \right\}, \end{aligned}$$

where $\lambda^* = \lambda |\mathcal{E}_k|^{1/p_k}$ and $\omega_1, \omega_2, \ldots, \omega_{p_k}$ are the ordered eigenvalues of Ω. Taking derivatives with respect to ω_j and setting equal to 0, we have:

$$\begin{aligned} & 1 - \frac{K}{\omega_j} + \frac{1}{p_k} \omega_j^{1/p_k - 1} \lambda^* \prod_{i \neq j} \omega_i^{1/p_k} = 0 \\ & \Leftrightarrow \omega_j = K - \frac{1}{p_k} \lambda^* \prod_{i=1}^{p_k} \omega_i^{1/p_k} \text{ for all } j = 1, \ldots, p_k. \end{aligned}$$

So all of the eigenvalues have the same critical value.

Taking second derivatives, we have:

$$\frac{K}{\omega_j^2} - \frac{p_k - 1}{p_k^2}\lambda^* \omega^{1/p_k - 2} \prod_{i \neq j}^{p_k} \omega_j^{1/p_k} > 0 \Leftrightarrow K - \frac{p_k - 1}{p_k^2}\lambda^* \prod_{j=1}^{p_k} \omega_j^{1/p_k} > 0$$

$$\Leftrightarrow K + \frac{p_k - 1}{p_k}\left(K - \frac{1}{p_k}\lambda^* \prod_{j=1}^{p_k} \omega_j^{1/p_k} - K\right) > 0$$

$$\Leftrightarrow K + \frac{p_k - 1}{p_k}(\omega_j - K) > 0 \Leftrightarrow \frac{p_k - 1}{p_k}\omega_j + K\frac{1}{p_k} > 0.$$

Hence, by a second derivative test, this critical value is a minimizer for all ω_j. This is a global minimum since

$$\text{as } \omega_1 \to \infty \text{ we have that } \left\{\sum_{i=1}^{p_k} \omega_i - K\sum_{i=1}^{p_k} \log(\omega_i) + \lambda^* \prod_{i=1}^{p_k} \omega_i^{1/p_k}\right\} \to \infty$$

and

$$\text{as } \omega_{p_k} \to 0 \text{ we have that } \left\{\sum_{i=1}^{p_k} \omega_i - K\sum_{i=1}^{p_k} \log(\omega_i) + \lambda^* \prod_{i=1}^{p_k} \omega_i^{1/p_k}\right\} \to \infty.$$

This implies that all of the ω_j are equal. In particular, that $\omega_j = (Kp_k)/(p_k + \lambda^*)$ for all $j = 1, \ldots, p_k$. This in turn implies that Ω is a constant multiple of the identity. Thus, the $\tilde{S}_k$ that minimizes the risk given all S_j such that $j \neq k$ is:

$$\tilde{S}_k = \frac{Kp_k}{p_k + \lambda^*}\mathcal{E}_k.$$

But this means that the S_k that minimizes this risk, no matter what the other S_j's are, is

$$\hat{\Sigma}_k = \mathcal{E}_k / |\mathcal{E}_k|^{1/(p_k)}.$$

It remains to minimize with respect to s. The minimizer is the s such that

$$2s\sum_{k=1}^{K} \frac{p}{p_k} \text{tr}\left(\hat{\Sigma}_k \mathcal{E}_k^{-1}\right) - \frac{2Kp}{s} = 0.$$

And solving for s we get

$$\hat{\sigma}^2 = \frac{K}{\sum_{k=1}^{K} \frac{1}{p_k} \text{tr}\left(\hat{\Sigma}_k \mathcal{E}_k^{-1}\right)}.$$

But since $\hat{\Sigma}_k = \mathcal{E}_k / |\mathcal{E}_k|^{1/(p_k)}$, we have that

$$\hat{\sigma}^2 = \frac{K}{\sum_{k=1}^{K} |\mathcal{E}_k|^{-1/p_k}}.$$

□

A.8 Proof of Proposition 3

Proof. Let $\Phi_k = \Sigma_k/\text{tr}(\Sigma_k)$, $D_k = S_k/\text{tr}(S_k)$ for $k = 1, \ldots, K$. So $\Sigma_k = \Phi_k/|\Phi_k|^{1/p_k}$ and $S_k = D_k/|D_k|^{1/p_k}$ for $k = 1, \ldots, K$. Φ_k and D_k both have trace 1. The space of trace 1 symmetric positive definite matrices is convex. Let $\Phi = (\sigma^2, \Phi_1, \ldots, \Phi_K)$ and $D = (s^2, D_1, \ldots, D_K)$. Define

$$L_2(\Phi, D) = \frac{s^2}{\sigma^2} \sum_{k=1}^{K} \frac{p}{p_k} |D_k \Phi_k^{-1}|^{-1/p_k} \text{tr}\left(D_k \Phi_k^{-1}\right) - Kp \log\left(\frac{s^2}{\sigma^2}\right) - Kp.$$

So, $L_M(\Sigma, S) = L_2(\Phi, D)$.

Hence, $E[L_M(\Sigma, S)|X] = E[L_2(\Phi, D)|X]$.

So if L_2 is convex in each D_k, we can uniformly decrease the risk. That is, given $B_k, E_k \in \mathcal{G}_{p_k}^+$ are two estimators from two different special linear group transformations, an estimator that uniformly decreases the risk is found by setting $F_k = (B_k/\text{tr}(B_k) + E_k/\text{tr}(E_k))/2$ and using $F_k/|F_k|^{1/p_k}$ as our estimator. Averaging over the whole space of orthogonal matrices will result in an orthogonally equivariant estimator.

It remains to prove that L_2 is convex in each D_k. It suffices to show that $|D_k|^{-1/p_k}\text{tr}(D_k\Phi_k^{-1})$ is convex in D_k. Since, for $\alpha \in [0,1]$, $\text{tr}((\alpha D_k + (1-\alpha)E_k)\Phi_k^{-1}) = \alpha\text{tr}(D_k\Phi_k^{-1}) + (1-\alpha)\text{tr}(E_k\Phi_k^{-1})$ is convex in D_k, if $|D_k|^{-1/p_k}$ is also convex, then we are done. $|D_k|$ is a concave function [Cover and Thomas, 1988, Theorem 1], and $f(x) = \log(x)$ is concave monotonic, so $\log(|D_k|)$ is concave, so $-\log(|D_k|)/p_k$ is convex, so $\exp(-\log(|D_k|)/p_k) = |D_k|^{-1/p_k}$ is convex.

We also have that $cb^2 - h\log(b^2)$ is convex in b^2 for $c, h > 0$, so we can average the scale estimates to decrease risk as well.

To summarize, we have:

$$\begin{aligned}
&L_M\left(\Sigma, \left(f^2, F_1/|F_1|^{1/p_1}, \ldots, F_K/|F_K|^{1/p_K}\right)\right) \\
&= L_2\left(\Phi, \left(f^2, F_1, \ldots, F_K\right)\right) \\
&= L_2\left(\Phi, \left((b^2+e^2)/2, B_1/\text{tr}(B_1) + E_1/\text{tr}(E_1)\right)/2,\right. \\
&\quad \left. \ldots, (B_K/\text{tr}(B_K) + E_K/\text{tr}(E_K))/2\right) \\
&\leq \frac{1}{2} L_2\left(\Phi, \left(b^2, B_1/\text{tr}(B_1), \ldots, B_K/\text{tr}(B_K)\right)\right) \\
&\quad + \frac{1}{2} L_2\left(\Phi, \left(e^2, E_1/\text{tr}(E_1), \ldots, E_K/\text{tr}(E_K)\right)\right) \\
&= \frac{1}{2} L_M(\Sigma, B) + \frac{1}{2} L_M(\Sigma, E).
\end{aligned}$$

If B and E have the same (constant) risk as the UMREE, $\hat{\Sigma}(X)$, then

$$\begin{aligned}
&E\left[L_M\left(\Sigma, \left(f^2, F_1/|F_1|^{1/p_1}, \ldots, F_K/|F_K|^{1/p_K}\right)\right)\right] \\
&\leq \frac{1}{2}E\left[L_M\left(\Sigma, B\right)\right] + \frac{1}{2}\left[L_M\left(\Sigma, E\right)\right] \\
&= E\left[L_M\left(\Sigma, \hat{\Sigma}(X)\right)\right]
\end{aligned}$$

□

References

Deniz Akdemir and Arjun K Gupta. Array variate random variables with multiway Kronecker delta covariance matrix structure. *Journal of algebraic statistics*, 2(1), 2011.

MS Bartlett. On the theory of statistical regression. *Proceedings of the Royal Society of Edinburgh*, 53:260–283, 1933.

James V Bondar and Paul Milnes. Amenability: A survey for statistical applications of Hunt-Stein and related conditions on groups. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 57(1):103–128, 1981.

Rasmus Bro. Review on multiway analysis in chemistry - 2000–2005. *Critical reviews in analytical chemistry*, 36(3-4):279–293, 2006.

A Cichocki, D Mandic, C Caiafa, AH Phan, G Zhou, Q Zhao, and L De Lathauwer. Tensor decompositions for signal processing applications. *From Two-way to Multiway Component Analysis, ESAT-STADIUS Internal Report*, pages 13–235, 2014.

Thomas M Cover and A Thomas. Determinant inequalities via information theory. *SIAM journal on Matrix Analysis and Applications*, 9(3):384–392, 1988.

A Philip Dawid. Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika*, 68(1):265–274, 1981.

Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. On the best rank-1 and rank-$(R_1, R_2, \ldots, R_n)$ approximation of higher-order tensors. *SIAM Journal on Matrix Analysis and Applications*, 21(4):1324–1342, 2000a.

Lieven De Lathauwer, Bart De Moor, and Joos Vandewalle. A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications*, 21(4):1253–1278, 2000b.

Vin De Silva and Lek-Heng Lim. Tensor rank and the ill-posedness of the best low-rank approximation problem. *SIAM Journal on Matrix Analysis and Applications*, 30(3):1084–1127, 2008.

Pierre Dutilleul. The MLE algorithm for the matrix normal distribution. *Journal of statistical computation and simulation*, 64(2):105–123, 1999.

Morris L Eaton. *Multivariate statistics: a vector space approach*. Wiley New York, 1983.

Morris L Eaton. Group invariance applications in statistics. In *Regional conference series in Probability and Statistics*, pages i–133. JSTOR, 1989.

Morris L Eaton, Ingram Olkin, et al. Best equivariant estimators of a cholesky decomposition. *The Annals of Statistics*, 15(4):1639–1650, 1987.

Montserrat Fuentes. Testing for separability of spatial–temporal covariance functions. *Journal of statistical planning and inference*, 136(2):447–466, 2006.

LR Haff. The variational form of certain Bayes estimators. *The Annals of Statistics*, pages 1163–1190, 1991.

Peter D Hoff. Separable covariance arrays via the Tucker product, with applications to multivariate relational data. *Bayesian Analysis*, 6(2):179–196, 2011.

Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.

William James and Charles Stein. Estimation with quadratic loss. In *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, volume 1, pages 361–379, 1961.

J Kiefer. Invariance, minimax sequential estimation, and continuous time processes. *The Annals of Mathematical Statistics*, 28(3):573–601, 1957.

Henk AL Kiers and Iven Van Mechelen. Three-way component analysis: Principles and illustrative application. *Psychological methods*, 6(1):84, 2001.

Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51 (3):455–500, 2009.

Samuel Kotz and Saralees Nadarajah. *Multivariate t-distributions and their applications*. Cambridge University Press, 2004.

Pieter M Kroonenberg. *Applied multiway data analysis*, volume 702. John Wiley & Sons, 2008.

SP Lin and Michael D Perlman. A Monte Carlo comparison of four estimators of a covariance matrix. *Multivariate Analysis*, 6:411–429, 1985.

X Magnus and Heinz Neudecker. Matrix differential calculus. *New York*, 1988.

Kanti V. Mardia. Spatial-temporal analysis of multivariate environmental monitoring data. In Ganapati P Patil and Calyampudi Radhakrishna Rao, editors, *Multivariate environmental statistics"*, volume 6. Elsevier, 1993.

Martin Ohlson, M Rauf Ahmad, and Dietrich Von Rosen. The multilinear normal distribution: Introduction and some basic properties. *Journal of Multivariate Analysis*, 113:37–47, 2013.

Joseph J Rotman. *An introduction to the theory of groups*, volume 148. Springer, 1995.

Mahendran Shitan and Peter J Brockwell. An asymptotic test for separability of a spatial autoregressive model. *Communications in Statistics-Theory and Methods*, 24(8):2027–2040, 1995.

Age Smilde, Rasmus Bro, and Paul Geladi. *Multi-way analysis: applications in the chemical sciences*. John Wiley & Sons, 2005.

MS Srivastava and CG Khatri. An introduction to multivariate statistics., 1979.

Charles Stein. Estimation of a covariance matrix. *Rietz Lecture*, 1975.

Akimichi Takemura. An orthogonally invariant minimax estimator of the covariance matrix of a multivariate normal population. Technical report, DTIC Document, 1983.

Dacheng Tao, Xuelong Li, Weiming Hu, Stephen Maybank, and Xindong Wu. Supervised tensor learning. In *Data Mining, Fifth IEEE International Conference on*, pages 8–pp. IEEE, 2005.

M Alex O Vasilescu and Demetri Terzopoulos. Multilinear subspace analysis of image ensembles. In *Computer Vision and Pattern Recognition, 2003. Proceedings. 2003 IEEE Computer Society Conference on*, volume 2, pages II–93. IEEE, 2003.

Ami Wiesel. On the convexity in Kronecker structured covariance estimation. In *Statistical Signal Processing Workshop (SSP), 2012 IEEE*, pages 880–883. IEEE, 2012a.

Ami Wiesel. Geodesic convexity and covariance estimation. *Signal Processing, IEEE Transactions on*, 60(12):6182–6189, 2012b.

Ruo-yong Yang and James O. Berger. Estimation of a covariance matrix using the reference prior. *Ann. Statist.*, 22(3):1195–1211, 1994. ISSN 0090-5364. doi: 10.1214/aos/1176325625. URL `http://dx.doi.org/10.1214/aos/1176325625`.

James V Zidek. A representation of Bayes invariant procedures in terms of Haar measure. *Annals of the Institute of Statistical Mathematics*, 21(1):291–308, 1969.